

Chapter 2

Defense in Depth

2.1 Introduction and Context Diagrams

Defense in Depth is a practical strategy for achieving information assurance (IA) in today's highly networked environments. It is a practical strategy because it relies on the intelligent application of techniques and technologies that exist today. This strategy recommends a balance among protection capability, cost, performance, and operational considerations. This chapter presents an overview of the major elements of this strategy and provides links to resources that offer additional insight.

2.1.1 Examples of User Environments

The following subsections introduce examples of customer computing environments and depict how they can interconnect with other organizational enclaves. The Information Assurance Technology Framework (IATF) technologies and suggested solutions provided apply to the computing environments described in these subsections. The Defense in Depth strategy and objectives described below apply equally to the federal computing environment and the Department of Defense (DoD) and concepts concerning computing environments. Defense of the computing environment, the enclave, and the network and infrastructure, apply in each environment in which all systems are interconnected.

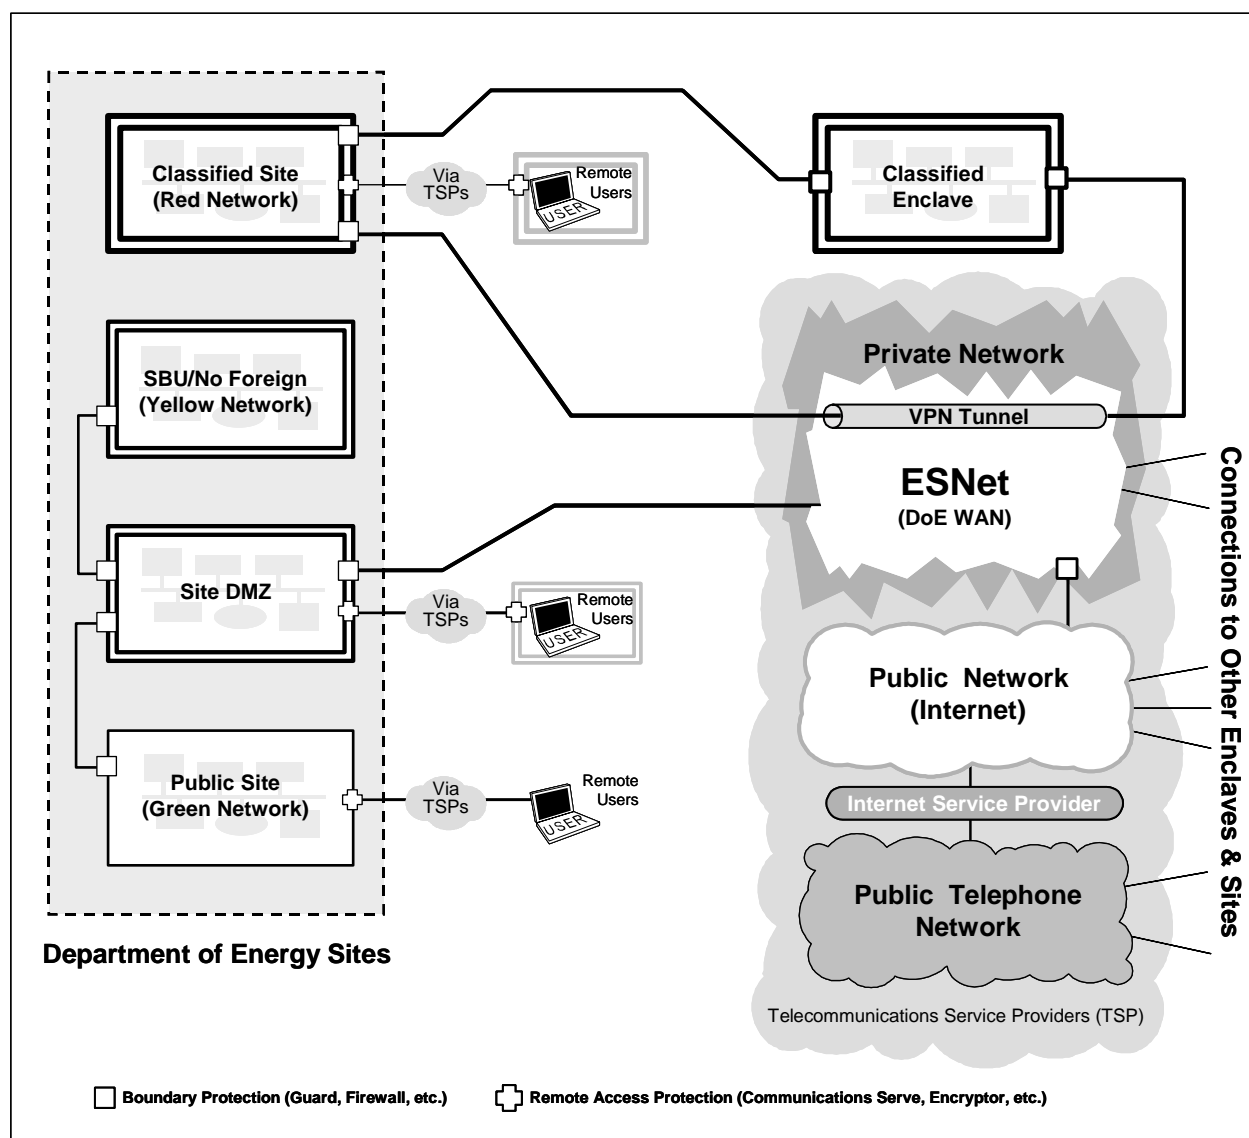
2.1.1.1 Federal Computing Environment

The interconnection of Department of Energy (DOE) research facilities, weapons laboratories, regional operations offices, and academic facilities is one example of a federal computing environment. The DOE information infrastructure is interconnected via several DOE wide area networks (WAN), one of which is the Energy Science Network (ESNet).

ESNet is a high-performance data communications backbone that provides DOE with widespread support for research and mission-critical applications. It supports both classified and unclassified DOE mission-oriented networking for scientists, engineers, and their administrative support. The ESNet consists of an asynchronous transfer mode (ATM) backbone and multiple local area networks (LAN) interconnected to establish a global network capability. ESNet permits virtual network architectures so that virtual networks can be layered on top of the existing network while running totally independent on the host network (i.e., ESNet). One DOE virtual network hosted on ESNet is SecureNet, a classified DOE support network. The virtual private network (VPN), SecureNet, provides a connection between three application-specific integrated circuits (ASIC) teraflop supercomputers, DOE headquarters, and other defense

program facilities across the United States. As a result, scientists and researchers at any of these DOE sites have on-demand access to the supercomputers.

Figure 2-1 presents a conceptual diagram of a typical DOE site within the broader DOE computing environment. The typical DOE site has two primary networks (or three, if the site processes classified information).



latf_2_2_1_0130

Figure 2-1. Federal Computing Environment—DOE

The primary networks include a “Green” unclassified or public network; a “Yellow” sensitive but unclassified/no-foreign (Unclassified but Controlled/NOFORN) network; and a “Red,” or classified, network. The Green, Yellow, and Red networks may each consist of one LAN or of multiple subnetworks. The typical DOE site has implemented a demilitarized zone (DMZ) or information protection network (IPN) that acts as the single point of entry into the site and

defends the enclave boundary or external connection(s). Within the Yellow and Red LANs, virtual networks are established to support various mission functions within the site. Physical isolation is primarily used to maintain the confidentiality and integrity of classified data. Carefully controlled connectivity is provided between the Red network, the Yellow network, and ESNet when data transfer outside the enclave is required.

All public information, Web-serve, and nonsensitive information are located on the Green network, which is normally protected by the site's DMZ resources. Remote access to the site will be established via the DMZ. A typical DOE site obtains Internet access via the ESNet connection.

2.1.1.2 DoD Computing Environment

The Defense Information Infrastructure (DII) environment is an example of one of the U.S. Government's largest and most complex information infrastructures. The DII supports more than 2 million primary users (with extensions to an additional 2 million users). Included within the DII are some 200 command centers and 16 large data centers, the Defense Megadata Centers. The basic user environments are enclaves (physically protected facilities and compounds), incorporating more than 20,000 local networks and some 4,000 connections to a backbone network. The DII also supports more than 300,000 secure telephone users.

The DII implements a number of global virtual networks that support a range of mission functions, for example, logistics, intelligence, and using WANs such as the Joint Worldwide Intelligence Communications System (JWICS) and the Secret Internet Protocol Router Network (SIPRNet) for global connectivity. In the past, this information infrastructure was based on dedicated networks and customized information systems; today, DoD is almost totally dependent on commercial services within the Nationwide Information Infrastructure (NII) and the broader global information infrastructure.

Figure 2-2 presents a system context diagram of a typical user site or facility within the broader DII structure. The typical user facility has several LANs that support the mission functional areas. Today, physical isolation is primarily used to maintain the confidentiality and the integrity of different classification levels of traffic. Within these isolated LANs, virtual networks are established to support the various mission functions within the enclave. Carefully controlled connectivity is provided between networks of different classification levels when boundaries are required.

For example, DoD organizations have robust, worldwide intelligence systems operating at top secret-sensitive compartmented information (TS-SCI) that carry significant levels of unclassified traffic. This supports the organizations' need to communicate with others within the intelligence community. Within the same TS-SCI enclaves, customers have secret and unclassified systems with less-than-robust connectivity to non-intelligence-community users. To reach a mixed community of users, unclassified information may have to flow over separate unclassified, secret, and TS-SCI systems. Moving information between these systems (enclaves) is complicated because of the need to comply with policy regarding releasability.

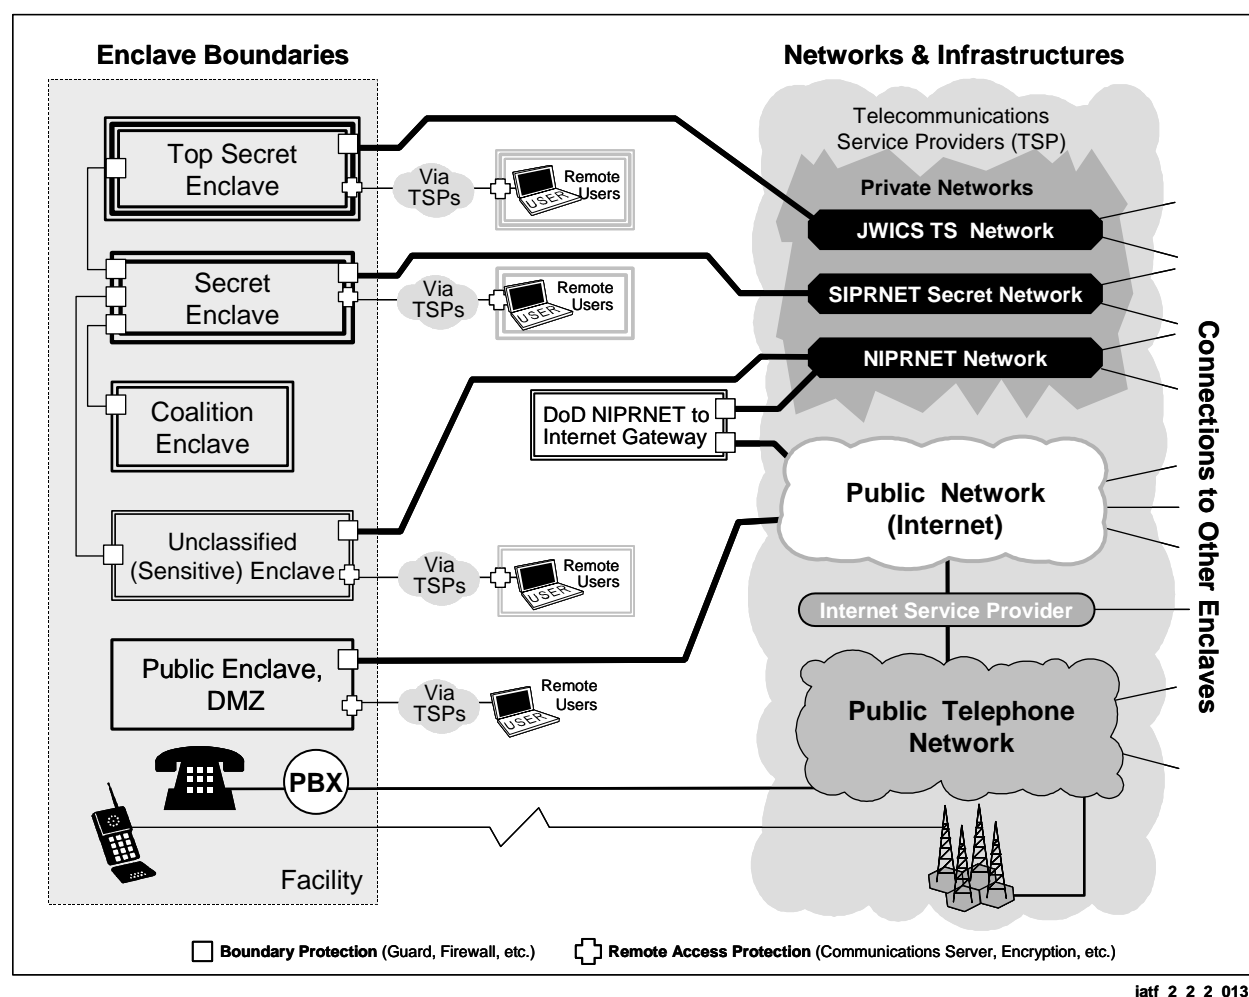


Figure 2-2. Federal Computing Environment—DoD

2.2 Adversaries, Motivations, and Classes of Attack

To effectively resist attacks on its information and information systems, an organization must characterize its adversaries, their potential motivations, and their attack capabilities. Potential adversaries might include nation states, terrorists, criminal elements, hackers, or corporate competitors. Their motivations might include intelligence gathering, theft of intellectual property, causing embarrassment, or just anticipated pride in having exploited a notable target. The methods of attack might include passive monitoring of communications, active network attacks, close-in attacks, exploitation of insiders, and attacks through the industry providers of the organization's information technology (IT) resources.

In addition to guarding against intentional attack, the organization must protect against the detrimental effects of nonmalicious events such as fire, flood, power outages, and user error.

For clarity and understanding, the rest of Section 2.2 is a reprint of Section 1.3 of the IATF.

Information systems and networks offer attractive targets. Therefore, they should be resistant to attack from the full range of threat agents—from hackers to nation states—and must be able to limit damage and recover rapidly when attacks do occur.

The IATF considers five classes of attacks:

- Passive.
- Active.
- Close-In.
- Insider.
- Distribution.

The key aspects of each class of attack are summarized in Table 2-1.

Table 2-1. Classes of Attack

Attack	Description
Passive	Passive attacks include traffic analysis, monitoring of unprotected communications, decrypting weakly encrypted traffic, and capture of authentication information (e.g., passwords). Passive intercept of network operations can give adversaries indications and warnings of impending actions. Passive attacks can result in disclosure of information or data files to an attacker without the consent or knowledge of the user. Examples include the disclosure of personal information such as credit card numbers and medical files.
Active	Active attacks include attempts to circumvent or break protection features, introduce malicious code, or steal or modify information. These attacks may be mounted against a network backbone, exploit information in transit, electronically penetrate an enclave, or attack an authorized remote user during an attempt to connect to an enclave. Active attacks can result in the disclosure or dissemination of data files, denial of service, or modification of data.
Close-In	Close-in attack consists of a regular type individuals attaining close physical proximity to networks, systems, or facilities for the purpose of modifying, gathering, or denying access to information. Close physical proximity is achieved through surreptitious entry, open access, or both.
Insider	Insider attacks can be malicious or nonmalicious. Malicious insiders intentionally eavesdrop, steal or damage information, use information in a fraudulent manner, or deny access to other authorized users. Nonmalicious attacks typically result from carelessness, lack of knowledge, or intentional circumvention of security for such reasons as “getting the job done.”
Distribution	Distribution attacks focus on the malicious modification of hardware or software at the factory or during distribution. These attacks can introduce malicious code into a product, such as a back door to gain unauthorized access to information or a system function at a later date.

The relationship of these attack classes to the information infrastructure areas is shown in Figure 2-3. Later sections of the IATF will provide an overview of the IA strategy for countering or mitigating the effects of these attacks.

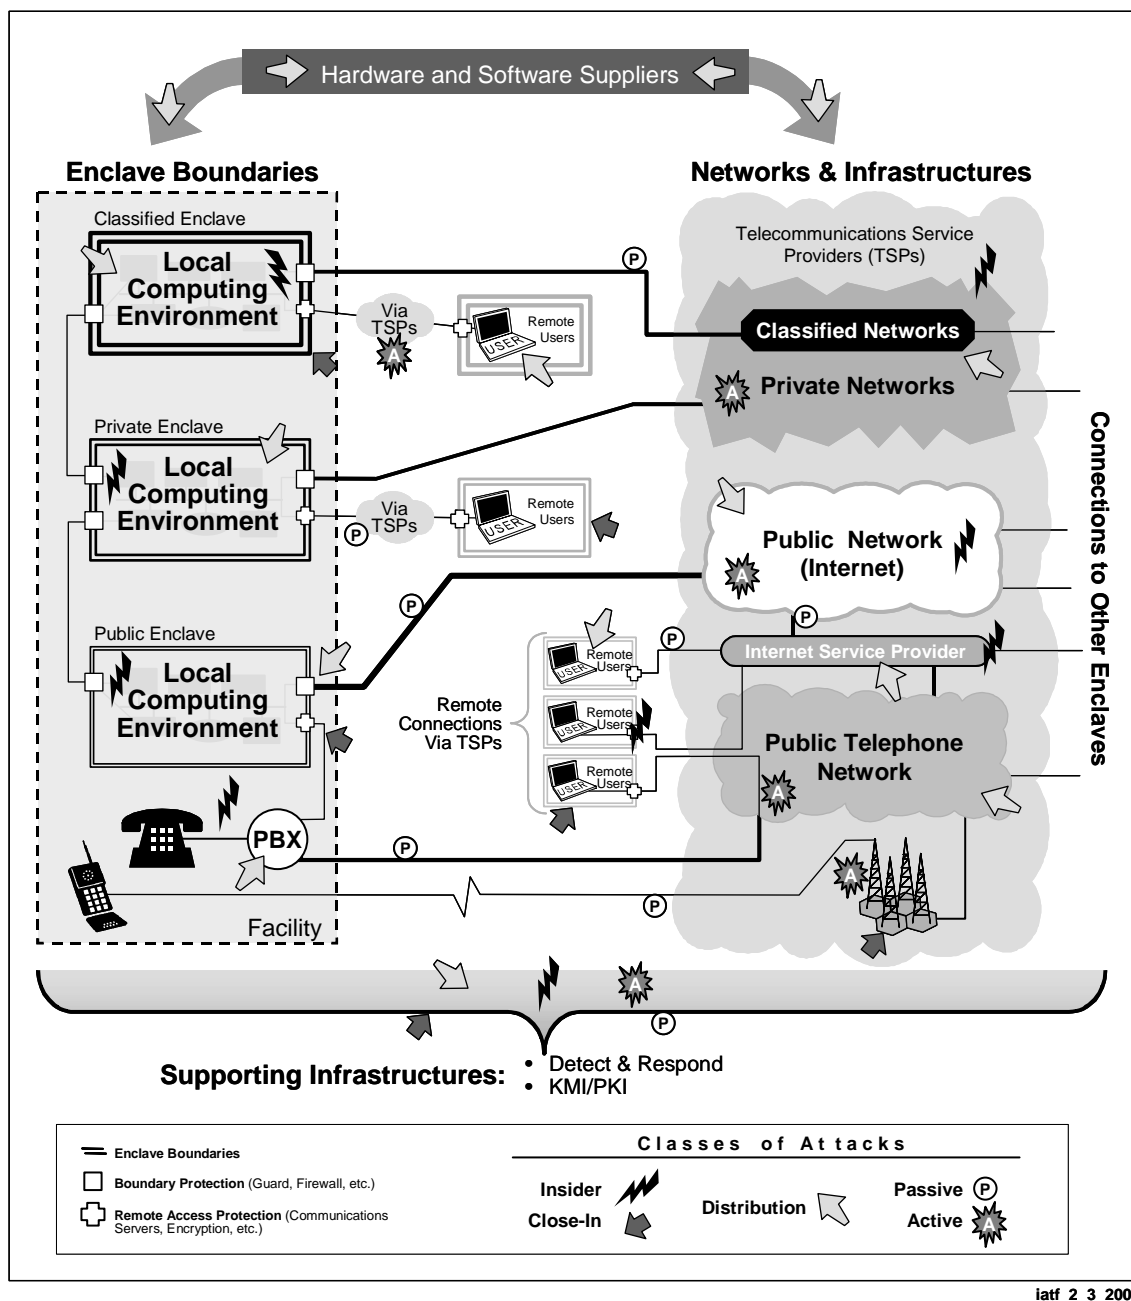


Figure 2-3. Classes of Attacks on the Information Infrastructure

2.3 People, Technology, Operations

IA is achieved when there is confidence that information and information systems are protected against attacks through the application of security services in such areas as availability, integrity, authentication, confidentiality, and nonrepudiation. The application of these services should be based on the protect, detect, and react paradigm. This means that in addition to incorporating protection mechanisms, organizations must expect attacks and must also incorporate attack-detection tools and procedures that allow them to react to and recover from these attacks.

Figure 2-4 depicts an important principle of the Defense in Depth strategy: the achievement of IA requires a balanced focus on three primary elements—people, technology, and operations.

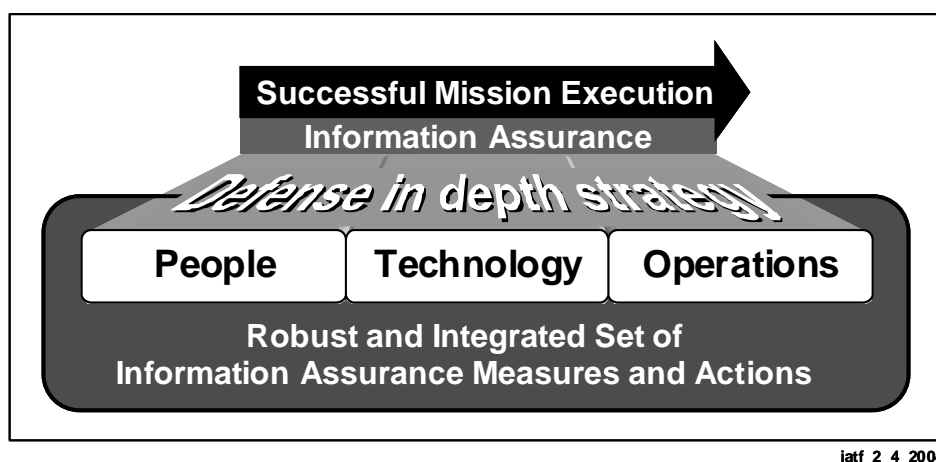


Figure 2-4. Defense in Depth Strategy

2.3.1 People

The achievement of IA begins with a senior-level management commitment (typically at the chief information officer level) based on a clear understanding of the perceived threat. This commitment must be followed by establishment of effective IA policies and procedures, assignment of roles and responsibilities, commitment of resources, training of critical personnel (e.g., users and system administrators), and enforcement of personal accountability. These steps include the establishment of physical security and personnel security measures to control and monitor access to facilities and critical elements of the IT environment.

Figure 2-5 lists some of the disciplines associated with people in the Defense in Depth strategy.

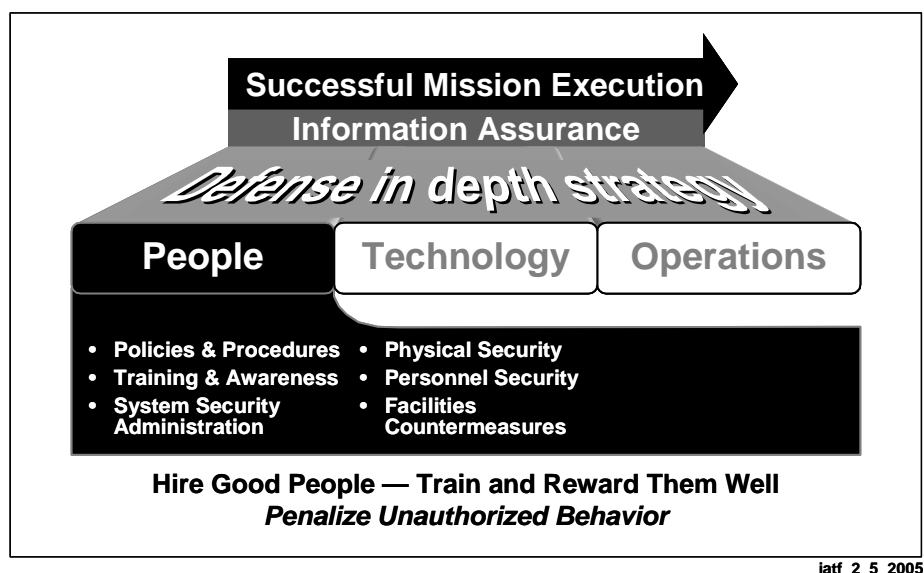


Figure 2-5. Defense in Depth Strategy—People

2.3.2 Technology

A wide range of technologies are available for providing IA services and for detecting intrusions. To ensure that the right technologies are procured and deployed, an organization should establish effective policies and processes for technology acquisition. These policies and processes should include security policy, IA principles, system-level IA architectures and standards, criteria for needed IA products, acquisition of products that have been validated by a reputable third party, configuration guidance, and processes for assessing the risk of the integrated systems. Figure 2-6 lists some of the technology areas addressed in the Defense in Depth strategy.

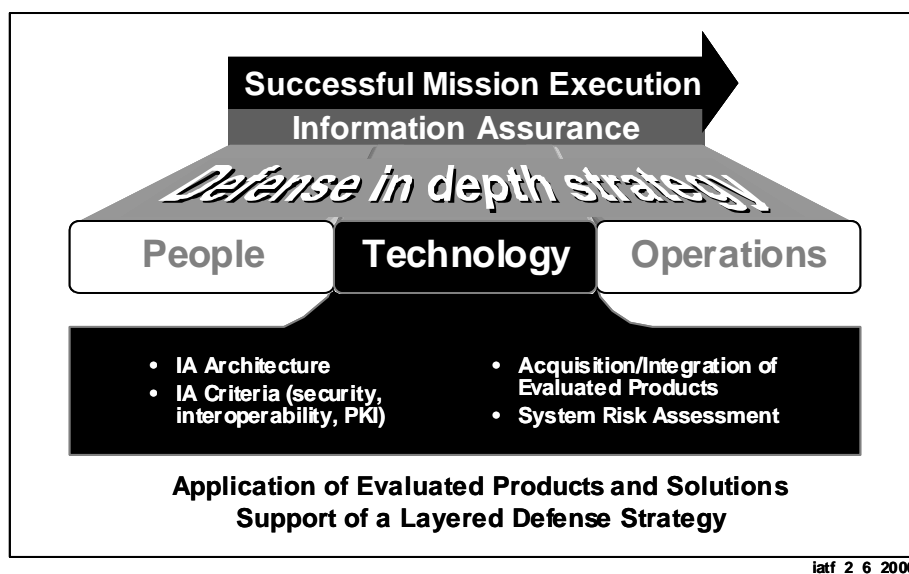
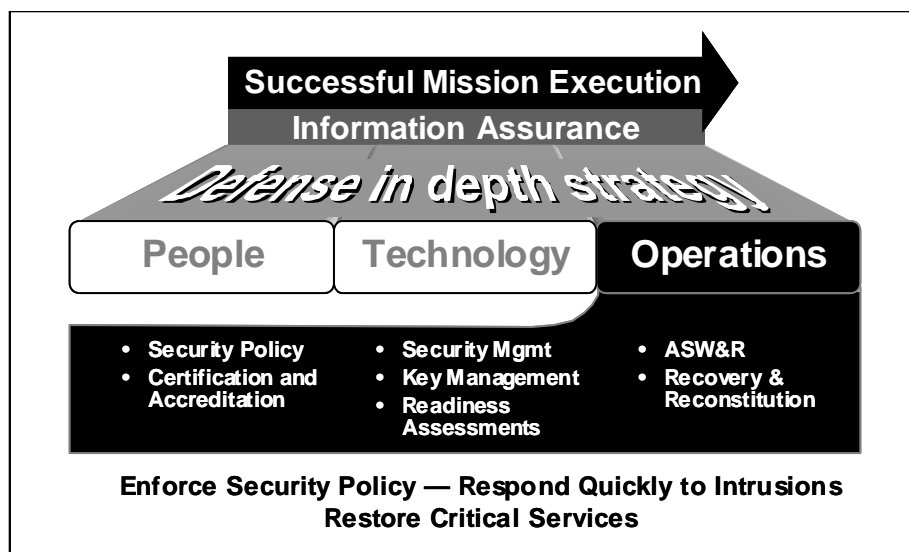


Figure 2-6. Defense in Depth Strategy—Technology

2.3.3 Operations

The operations element of the strategy focuses on all activities required to sustain an organization's security posture on a day-to-day basis. Figure 2-7 lists some of the operational focus areas associated with the Defense in Depth strategy.



iatf_2_7_2007

Figure 2-7. Defense in Depth Strategy—Operations

2.4 Defense in Depth Objectives Overview

The need for secure operations of information and communications systems is not new. However, as organizations' reliance on such systems increases, as entities strive for greater efficiency through shared resources, and as those who perpetrate threats become more numerous and more capable, the IA posture of systems and organizations grows ever more important. Deliberate investments of time, resources, and attention in implementing and maintaining an effective IA posture have never been more important or more challenging.

In implementing an effective and enduring IA capability or in adopting a Defense in Depth strategy for IA, organizations should consider—

- Taking into consideration the effectiveness of the information protection required, based on the value of the information to the organization and the potential impact that loss or compromise of the information would have on the organization's mission or business. IA decisions should be based on risk analysis and keyed to the organization's operational objectives.

UNCLASSIFIED

Defense-in-Depth
IATF Release 3.1—September 2002

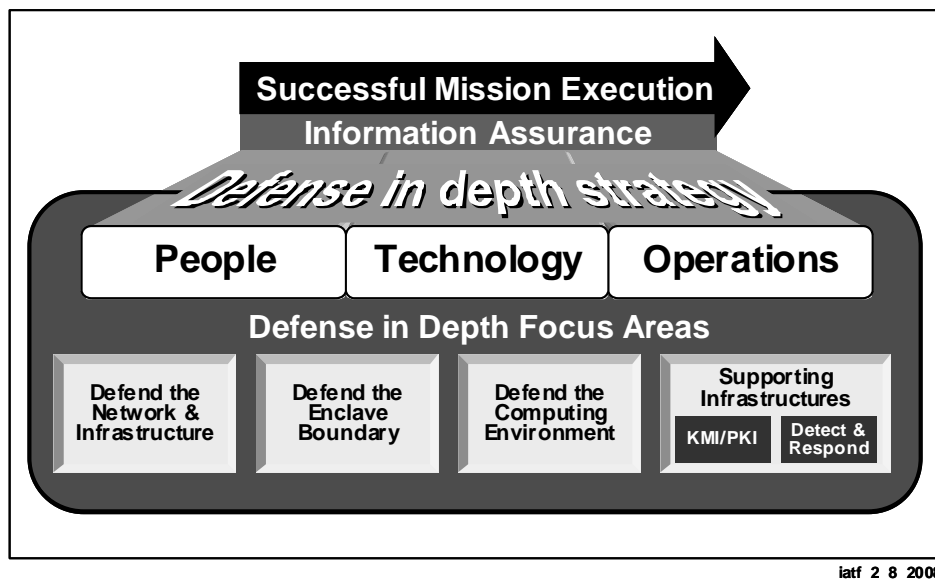
- Using a composite approach, based on balancing protection capability against cost, performance, operational impact, and changes to the operation itself considering both today's and tomorrow's operations and environments.
- Drawing from all three facets of Defense in Depth—people, operations, and technology. Technical mitigations are of no value without trained people to use them and operational procedures to guide their application.
- Establishing a comprehensive program of education, training, practical experience, and awareness. Professionalization and certification licensing provide a validated and recognized expert cadre of system administrators.
- Exploiting available commercial off-the-shelf (COTS) products and relying on in-house development for those items not otherwise available.
- Planning and following a continuous migration approach to take advantage of evolving information processing and network capabilities—both functional and security-related—and to ensure adaptability to changing organizational needs and operating environments.
- Assessing periodically the IA posture of the information infrastructure. Technology tools, such as automated scanners for networks, can assist in vulnerability assessments.
- Taking into account, not only the actions of those with hostile intent, but also inadvertent or unwitting actions that may have ill effects and natural events that may affect the system.
- Adhering to the principles of commonality, standardization, and procedures, and interoperability and to policies.
- Judiciously using emerging technologies, balancing enhanced capability with increased risk.
- Employing multiple means of threat mitigation, overlapping protection approaches to counter anticipated events so that loss or failure of a single barrier does not compromise the overall information infrastructure.
- Implementing and holding to a robust IA posture—one that can cope with the unexpected.
- Ensuring that only trustworthy personnel have physical access to the system. Methods of providing such assurance include appropriate background investigations, security clearances, credentials, and badges.
- Monitoring vulnerability listings and implementing fixes, ensuring that security mechanisms are interoperable, keeping constant watch over the security situation and mechanisms, properly employing and upgrading tools and techniques, and dealing rapidly and effectively with issues.

- Using established procedures to report incident information provided by intrusion detection mechanisms to authorities and specialized analysis and response centers.

The dominant need of the user community is ready access to the information infrastructure and the information it contains to support its operational objectives. This requires the use of robust information-processing technology and reliable connectivity. IA enables these capabilities by providing organizations with the ability to maintain adequate protection of their information.

The IATF focuses on the technology aspects of Defense in Depth. In developing an effective IA posture, all three components of the Defense in Depth strategy—people, technology, and operations—must be addressed.

The IATF organizes the presentation of IA technology objectives and approaches for the information infrastructure according to the four Defense in Depth technology focus areas: defend the computing environment, defend the enclave boundaries, defend the networks and infrastructure, and supporting infrastructures. These areas are shown in Figure 2-8. The technology objectives and approaches in these focus areas, explained in the sections that follow, address the needs of private and public, as well as civil and military, sectors of our society.



iatf_2_8_2008

Figure 2-8. Defense in Depth Focus Areas

The Defense in Depth strategy recommends adherence to several IA principles—

- **Defense in Multiple Places.** Given that adversaries can attack a target from multiple points using insiders or outsiders, an organization must deploy protection mechanisms at multiple locations to resist all methods of attack.

At a minimum, these Defense in Depth locations should include—

- Defend the networks and infrastructure:
 - Protect local and wide area communications networks (e.g., from denial-of-service attacks).
 - Provide confidentiality and integrity protection for data transmitted over these networks (e.g., use encryption and traffic flow security measures to resist passive monitoring).
 - Ensure that all data exchanged over WAN is protected from disclosure to anyone not authorized to access the network.
 - Ensure that WANs supporting mission-critical and mission-support data provide appropriate protection against denial-of-service attacks.
 - Protect against the delay, misdelivery, or nondelivery of otherwise adequately protected information.
 - Protect from traffic flow analysis:
 - User traffic.
 - Network infrastructure control information.
 - Ensure that protection mechanisms do not interfere with otherwise seamless operation with other authorized backbone and enclave networks.
- Defend the enclave boundaries (e.g., deploy firewalls and intrusion detection to resist active network attacks).
 - Ensure that physical and logical enclaves are adequately protected.
 - Enable dynamic throttling of services in response to changing threats.
 - Ensure that systems and networks within protected enclaves maintain acceptable availability and are adequately defended against denial-of-service intrusions.
 - Ensure that data exchanged between enclaves or via remote access is protected from improper disclosure.
 - Provide boundary defenses for those systems within the enclave that cannot defend themselves due to technical or configuration problems.
 - Provide a risk-managed means of selectively allowing essential information to flow across the enclave boundary.
 - Provide protection against the undermining of systems and data within the protected enclave by external systems or forces.
 - Provide strong authentication, and thereby authenticated access control, of users sending or receiving information from outside their enclave.
- Defend the computing environment (e.g., provide access controls on hosts and servers to resist insider, close-in, and distribution attacks).
 - Ensure that clients, servers, and applications are adequately defended against denial of service, unauthorized disclosure, and modification of data.
 - Ensure the confidentiality and integrity of data processed by the client, server, or application, both inside and outside of the enclave.
 - Defend against the unauthorized use of a client, server, or application.
 - Ensure that clients and servers follow secure configuration guidelines and have all appropriate patches applied.

- Maintain configuration management of all clients and servers to track patches and system configuration changes.
 - Ensure that a variety of applications can be readily integrated with no reduction in security.
 - Ensure adequate defenses against subversive acts by trusted persons and systems, both internal and external.
- **Layered defenses.** Even the best available IA products have inherent weaknesses. As a result, an adversary will eventually find an exploitable vulnerability in almost any system. An effective countermeasure is to deploy multiple defense mechanisms between adversaries and their target. Each of these mechanisms must present unique obstacles to the adversary. Further, each should include both protection and detection measures. These measures help to increase risk (of detection) for the adversary while reducing his or her chances of success or making successful penetrations unaffordable. Deploying nested firewalls (each coupled with intrusion detection) at outer and inner network boundaries is an example of a layered defense. The inner firewalls may support more granular access control and data filtering. Table 2-2 provides other examples of layered defenses.

Table 2-2. Examples of Layered Defenses

Class of Attack	First Line of Defense	Second Line of Defense
Passive	Link and network layer and encryption and traffic flow security	Security-enabled applications
Active	Defend the enclave boundaries	Defend the computing environment
Insider	Physical and personnel security	Authenticated access controls, audit
Close-In	Physical and personnel security	Technical surveillance countermeasures
Distribution	Trusted software development and distribution	Run time integrity controls

- **Security robustness.** Specify the security robustness (strength and assurance) of each IA component as a function of the value of what it is protecting and the threat at the point of application.
- **Deploy KMI/PKI.** Deploy robust key management and public key infrastructures that support all of the incorporated IA technologies and that are highly resistant to attack. Provide a cryptographic infrastructure that supports key, privilege, and certificate management and that enables positive identification of individuals using network services.
- **Deploy intrusion detection systems.** Deploy infrastructures to detect intrusions, to analyze and correlate the results, and to react as needed. These infrastructures should help the Operations staff to answer questions such as “Am I under attack?” “Who is the source?” “What is the target?” “Who else is under attack?” “What are my options?”

- Provide an intrusion detection, reporting, analysis, assessment, and response infrastructure that enables rapid detection and response to intrusions and other anomalous events and provides operational situation awareness.
- Plan execution and reporting requirements for contingencies and reconstitution.

2.5 Additional Resources

The National Security Agency (NSA), with support from other U.S. government agencies and U.S. industry, has undertaken several initiatives to support the Defense in Depth strategy. These include—

- **The IATF and the IATF Forum (www.iatf.net).** This document and the associated forum provide a means for the Government and industry to encourage a dialogue on IA issues.
- **The National Information Assurance Partnership (NIAP).** This is a partnership between NSA and the National Institute of Standards and Technology (NIST) to foster development of the International Common Criteria (an ISO standard) and to accredit commercial laboratories to validate the security functions in vendors' products. Information on this activity is available at <http://niap.nist.gov>.
- **Common Criteria Protection Profiles.** These documents recommend security functions and assurance levels based on the Common Criteria. They are available for a wide range of commercially available technologies and can be accessed at the IATF Web site (www.iatf.net) or the NIAP Web site (listed above).
- **List of Evaluated Products.** These are lists of commercial IA products that have been evaluated against the Common Criteria. The lists are maintained by NIST and are available at the NIAP Web site.
- **Configuration Guidance.** These documents, prepared by NSA, contain recommended configurations for a variety of commonly used commercial products. These documents can be found at <http://nsa1.www.conxion.com>.
- **Glossary.** *The National Information Systems Security (INFOSEC) Glossary* (September 2000) can be found at <http://www.nstissc.gov/Assets/pdf/4009.pdf>